Computationally Efficient Implementation of a Hamming Code Decoder using Graphics Processing Unit

Md Shohidul Islam, Cheol-Hong Kim, and Jong-Myon Kim*

arXiv:1412.6862v1 [cs.DC] 22 Dec 2014

Abstract: **This paper presents a computationally efficient implementation of a Hamming code decoder on a graphics processing unit (GPU) to support real-time software-defined radio (SDR), which is a software alternative for realizing wireless communication. The Hamming code algorithm is challenging to parallelize effectively on a GPU because it works on sparsely located data items with several conditional statements, leading to non-coalesced, long latency, global memory access, and huge thread divergence. To address these issues, we propose an optimized implementation of the Hamming code on the GPU to exploit the higher parallelism inherent in the algorithm. Experimental results using a compute unified device architecture (CUDA)-enabled NVIDIA GeForce GTX 560, including 335 cores, revealed that the proposed approach achieved a 99x speedup versus the equivalent CPU-based implementation.**

Index Terms: **Hamming code, GPU optimization, Software-defined radio.**

I. INTRODUCTION

Many existing wireless communication systems employed application specific integrated circuits (ASICs) based dedicated devices for particular communication protocol standards, including worldwide interoperability for microwave access (WiMAX, IEEE 802.16), Wi-Fi (IEEE802.11), digital high definition TV, wideband code division multiple access (W-CDMA), and global system for mobile communication (GSM) [1]-[7]. However, the fixed functionality of such ASIC devices limits their application to emerging communication standards because they were fixed for specific coding schemes, data rates, frequency ranges, and types of modulation [8]. In addition, manufacturing costs are high and time-to-market of hardware devices is long [9].

Software-defined radio (SDR) is an emerging technology that offers software alternatives to existing hardware solutions for wireless communication [10],[11], SDR technology has recently attracted the interest of the communication research community [12],[13]. SDR comprises software implementation of multi-standard and multi-protocol communication systems using one hardware platform [14]. It allows system reconfiguration by using software commands, because users are required to switch from one standard to another standard very frequently [15]. In addition, it enables the radio device to change transmitting and receiving characteristics by means of the software without altering the hardware platform [15]. In SDR, some or all of the physical layer functions are coded in the software, which runs on general-purpose programmable processors (GPPs) and digital signal processors (DSPs) [16],[17]. GPPs and DSPs offer the necessary programmability and flexibility for various SDR applications. However, neither GPPs nor DSPs can meet the much higher levels of performance required by high computational workloads in SDR [18].

Among many available computational models, graphics processing units (GPUs) perform well when performing latency-tolerant, highly parallel, and independent tasks. Attracted by the features of modern GPUs, many researchers have developed GPU-based SDR systems including turbo decoders, LDPC decoders, Viterbi decoders, and MIMO detectors to meet the high throughput required by the SDR algorithm [19]-[25]. In this paper, we present an optimized implementation of a Hamming decoder on a GPU; the Hamming decoder is widely used as a forward error correction (FEC) mechanism in wireless communication. Practical applications of the Hamming decoder include Ethernet (IEEE 802.3), WiMAX (IEEE 802.16e), Wi-Fi (IEEE 802.11n), telecommunication, digital video broadcasting-satellite second generation (DVB-S2), wireless sensor networks (WSNs), underwater wireless sensor networks (UWSNs), and space communication [26]-[31].

Contributions of this study are as follows:

- This paper presents a massively parallel and optimized implementation of a Hamming decoder on a GPU by exploring memory transfer, memory transaction, and kernel computation.
- The performance of the Hamming decoder on the GPU is thoroughly evaluated for various packet sizes, code lengths, and error tolerance.
- The performance of the proposed GPU approach is compared with the equivalent sequential approach run on a conventional CPU.

The remainder of this paper is organized as follows. A review of the Hamming decoder is provided in Section II, optimization and GPU implementation of the Hamming decoder are presented in Section III, and experimental results and analysis are discussed in Section IV.

Manuscript received Feb. 5, 2014: approved for publication by Prof. Ali Abedi, Dec. 1, 2014. This work was supported by a grant from the National Research Foundation of Korea (NRF) funded by the Korean government (MEST) (No. NRF-2013R1A2A2A05004566).

Md Shohidul Islam and Jong-Myon Kim are with the School of Electrical Engineering, University of Ulsan, Korea, email: shohid@mail.ulsan.ac.kr,jmkim07@ulsan.ac.kr. *Jong-Myon Kim is a corresponding author.

Cheol-Hong Kim is with the School of Electronics and Computer Engineering, Chonnam National University, Gwangju, Korea, email: chkim22@chonnam.ac.kr.

1229-2370/14/$10.00 © 2014 KICS

II. REVIEW OF THE HAMMING DECODER

Hamming decoding is performed at the destination end of the packet, and involves the exact reverse process of encoding performed at the transmitter end. Figure 1 shows a Hamming code decoder that consists of three components: splitter, decoder, and merger. The splitter receives the Hamming encoded packet, $\mathbf{H}=\{\mathbf{b}; \mathbf{b}=0 \mid \mathbf{b}=1\}$, and splits the message into t segments, $\mathbf{H}_1$, $\mathbf{H}_2$,..., $\mathbf{H}_t$, where *t* is the error tolerance. The main decoder consists of three fundamental units: error detection (ED), error correction (EC), and redundancy remover (RR). We use the terms *packet* and *message* interchangeably throughout this paper.

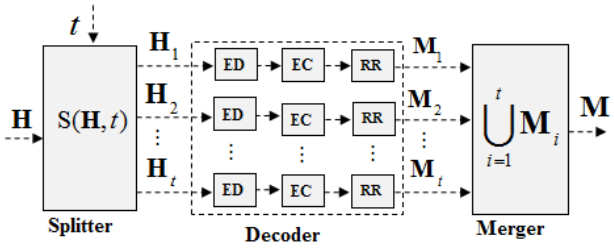

Fig. 1. Components in a Hamming code decoder.

In the encoding process, some redundancy or checksum bits are incorporated along with the original message for the purpose of error detection, and these should be removed once they have served their purpose. Subsequently, the decoder retrieves message segments $\mathbf{M}_1$, $\mathbf{M}_2$,..., $\mathbf{M}_t$, and the merger unifies them to produce the decoded packet, **M**, which is similar to the original packet sent by the sender or transmitter.

III. GPU-BASED IMPLEMENTATION OF THE HAMMING DECODER

This section presents a computationally efficient implementation of the Hamming code algorithm on a GPU. Encoded packets, namely P_1, P_2,...,P_n, are primarily received in the receiver buffer and the entire task can be divided into three steps, as shown in Figure 2: (i) pre-processing in the CPU, (ii) packet transfer between the CPU and GPU, and (iii) device kernel execution (DKE). All of these steps are performed in an optimized manner from a GPU computing viewpoint.

Figure 3 depicts an execution flow of the entire decoding process in the destination end of a network data packet, where regular blocks represent the steps executed on the CPU and the dotted blocks represent the tasks in the GPU. The CPU and GPU are also called the Host and Device, respectively. At the outset, the encoded packet, **H**, undergoes pre-processing in the CPU, which is explained in Section 3.1, before being transferred to the GPU. A parallel algorithm executed on the GPU is called a kernel, and the proposed approach configures two kernels, namely *checksum* and *error*, as indicated in Figure 3. The *checksum* kernel computes the redundancy information, and the *error* kernel performs error detection, correction, redundancy removal, and finally retrieves the original packet. This packet, now referred to as the decoded packet, is transferred from the GPU back to the CPU.

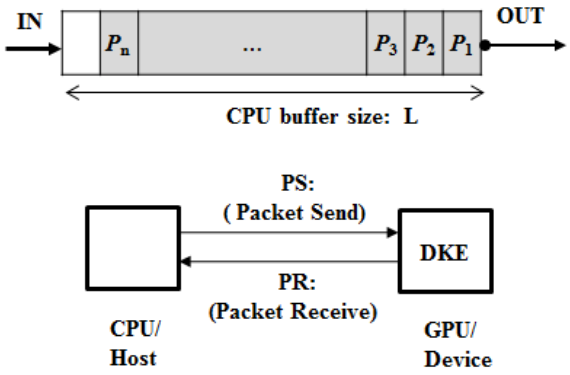

Fig. 2. Task partitioning for the proposed Hamming decoder.

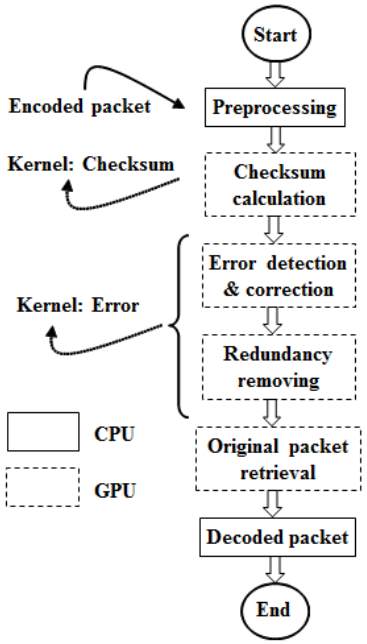

Fig. 3. Execution flow of the Hamming decoding procedure on a CPU and GPU.

A. Packet pre-processing

Instead of transferring the encoded packet, $\mathbf{H}=\mathbf{H}_1+\mathbf{H}_2+...+\mathbf{H}_t$, to the GPU immediately, packet pre-processing is first performed, because the first step in a GPU is to calculate checksums on each of $\mathbf{H}_1$,$\mathbf{H}_2$,...,$\mathbf{H}_t$; this process accesses sparsely located elements in global memory. Considering any $\mathbf{H}_i$, the index sets $\mathbf{I}_0$, $\mathbf{I}_1$,...,$\mathbf{I}_{|\mathbf{R}|-1}$, shown in Figure 4, access indices of $\mathbf{H}_i$ that are not completely adjacent. For instance, (7, 4) Hamming code has $|\mathbf{H}_i|=7+4=11$ bits and $|\mathbf{R}|=4$. Thus, $\mathbf{I}_0=\{1,3,5,7,9,11\}$, $\mathbf{I}_1=\{2,3,6,7,10,11\}$, $\mathbf{I}_2=\{4,5,6,7\}$, and $\mathbf{I}_3=\{8,9,10,11\}$. Consequently, data items accessed by $\mathbf{I}_0$ and $\mathbf{I}_1$are clearly non-adjacent. Even though elements of $\mathbf{I}_2$ and $\mathbf{I}_3$ apparently seem to be adjacent for a small code length, they include non-adjacent indices for larger values of $|\mathbf{H}_i|$.

If the packet segment is transferred to the GPU without pre-processing, there are two major performance bottlenecks: (a)

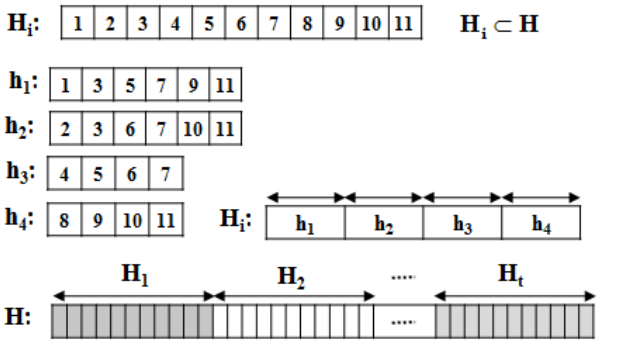

Fig. 4. Encoded packet re-construction for achieving coalesced global memory access in GPU.

non-adjacent memory transactions from GPU global memory result in long latency for memory READ, and (b) a number of conditional statements are required to access those locations, leading to thread divergence. These issues are addressed by pre-processing, which achieves coalesced global memory access. To this end, we re-organize the data items of the message segment, $\mathbf{H}_i$, by arranging those sparse items together and forming clusters such as $\mathbf{h}_1$, $\mathbf{h}_2$, $\mathbf{h}_3$, $\mathbf{h}_4$ for each group. The clusters are placed side-by-side to shape the new $\mathbf{H}_i$ as shown in Figure 4 Finally, the reformed encoded packet, $\mathbf{H}$, is created by concatenation of $\mathbf{H}_i{}'$s such that i=1,2,...,t, and this reformed encoded packet is transferred to the GPU.

B. Packet transfer between CPU and GPU

Data transfer between host and device is a vital issue in GPU computing. We utilize an optimized data transfer approach to achieve high performance. A GPU facilitates two modes of data transfer: synchronous data transfer (SDT) and asynchronous data transfer (ADT). Referring to Figure 2, three independent tasks, namely encoded packet transfer to the GPU [PS], DKE, and decoded packet receive from the GPU [PR], are accomplished in the GPU. As shown in Figure 5, these tasks are performed concurrently in ADT, where most of the transfer time is hidden by kernel execution. In contrast, SDT takes a long time and completes tasks in a sequential manner. As a result, ADT outperforms SDT by due to its pipelined execution pattern; therefore, we utilize ADT.

Figure 5 depicts the key differences between SDT and ADT in terms of execution time line. T_{PS}, T_{DKE}, and T_{PR} indicate the time required for the PS, DKE, and PR, respectively. The SDT based approach takes $3\times T_{PS}$ + $3\times T_{DKE}$+$3\times T_{PR}$ time units to process three packets, whereas ADT requires T_{PS} + $3\times T_{DKE}$ + T_{PR}. Consequently, ADT saves $2(T_{PS} + T_{PR})$ time units. In general, the speedup of ADT over SDT for the processing of N packets can be expressed by

$$ADT\ Speedup = \frac{SDT\ Execution\ Time}{ADT\ Execution\ Time} \qquad (1)$$

$$= \frac{N \times (T_{PS} + T_{DKE} + T_{PR})}{T_{PS} + N \times T_{DKE} + T_{PR}}$$

$$\approx (N, ..., 2); when(T_{PS}+T_{PR} \leq T_{DKE}, ..., T_{PS}+T_{PR} = T_{DKE})$$

Therefore, the minimum expected gain is two times; in our implementation, T_{PS} + $T_{PR} \geq T_{DKE}$, which accelerates the process toward N.

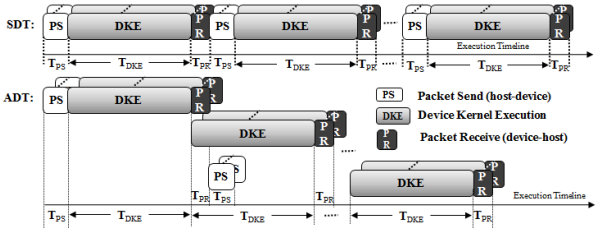

Fig. 5. Packet transfer using ADT and SDT.

C. Device Kernel Execution (DKE)

Algorithms that are executed in parallel on a GPU are called kernels. Algorithm 1 and Figure 6 show a *checksum* kernel that accesses the pre-processed encoded packet in the GPU global memory. The main task of this kernel is to calculate a checksum vector, $\mathbf{C}_i$, on $\mathbf{H}_i$, which is the major computation of the Hamming decoder.

Algorithm 1: $Kernel - Checksum(\mathbf{H}_1, \mathbf{C}_1)$

Input: Packet segment, $\mathbf{H}_1$, from the pre-processed encoded packet, $\mathbf{H}$
Output: Checksum vector, $\mathbf{C}_1$
Step 1: READ $\mathbf{H}_1$ from global memory
Step 2: For each block of GPU in parallel, do segmentation on $\mathbf{H}_1$ by index set $\mathbf{I}_0, \mathbf{I}_1, ..., \mathbf{I}_{r(\mathbf{H}_1)}$
Step 3: WRITE segments to the shared memory of each block
Step 4: Perform module 2 (XOR) operation on the shared memory packet segment
Step 5: WRITE the result of Step 4, the checksum vector $\mathbf{C}_1[1], \mathbf{C}_1[2], ..., \mathbf{C}_1[r(\mathbf{H}_1)]$, in global memory

The notation $r(\mathbf{H}_i)$ corresponds to the number of redundant or checksum bits required for error detection in the i^{th} message segment, $\mathbf{H}_i$. The kernel declares $r(\mathbf{H}_i)$ blocks, where each block is responsible for calculating one checksum bit. The checksum calculation is done by a modulo 2 operation on the bit strings stored in each block's shared memory by applying equations, as shown in Figure 6. The proposed approach removes bank conflicts by accessing items from a different shared memory bank. This kernel is accelerated by three optimizations: (i) coalesced global memory access, (ii) bank conflict avoidance by a reduction tree, and (iii) most computations based on shared memory.

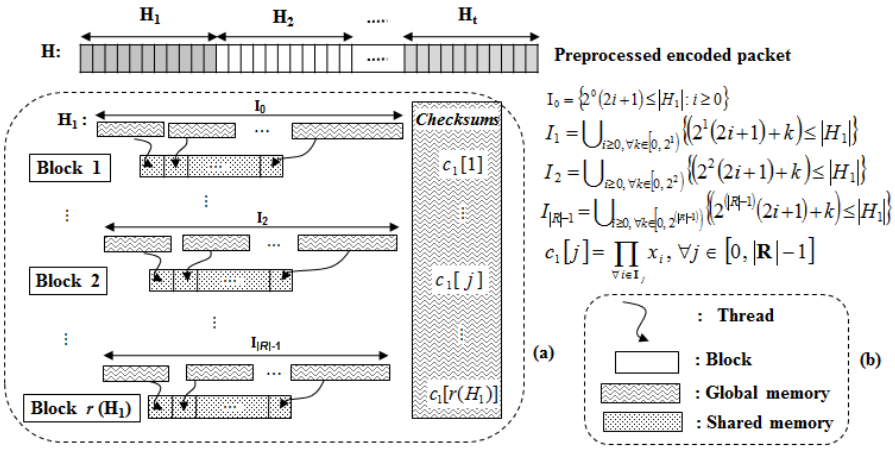

Fig. 6. Kernel *checksum* to calculate checksum on packet segment $\mathbf{H}_1$.

IV. Experimental Results and Analysis

To execute the CPU code, we use a machine running on Windows 7 (32 bit) with a 4-core 3.40GHz Intel processor that utilizes 8GB main memory. We evaluate the performance of the proposed GPU implementation on a 1.62GHz NVIDIA GeForce GTX 560 GPU with seven streaming multiprocessors (SM) and 1GB of main memory, where the GPU employed 336 processing elements and utilized 49,152 bytes of shared memory per SM. Furthermore, the maximum threads per block are 1,024 and the warp size is 32 threads. In this section, the execution times of the GPU-based and CPU-based Hamming decoders are compared.

A. Execution Time

Figures 7 and 8 show the consolidated execution time according to packet size, **M**, and error tolerance, *t*, for sequential and parallel Hamming decoding on the CPU and GPU, respectively. The packet length ranges from 400 bytes to 2000 bytes (x-axis) and the error tolerance is tested from 2 bits up to 6 bits (y-axis); the decoding time in milliseconds (*ms*) is shown on the z-axis.

Execution time generally increases with packet size for the following two reasons. First, longer messages have a greater number of checksum bits attached to the message. Second, the code used to calculate redundant information is also lengthened, increasing the number of XOR operations required. In addition, the computational time of CPU implementation is proportionally influenced by the Hamming code length of the packet. In contrast, the time for the GPU to decode packages remains relatively constant regardless of packet size.

There is a gradual increase in execution time as error tolerance increases for the CPU, while execution time of the GPU decreases as error tolerance, *t*, increases. The increase in *t* implies that a large number of bits in the packet are corrupted, because the transmission medium is erroneous. In these situations, the decoder splits away the packet into a higher number of segments, as shown in Figure 1. The CPU-based approach finishes these segments in a sequential manner, leading to an increase in decoding time. In contrast, the GPU deals with the segments in parallel and thus mitigates the increase in *t* with error tolerance.

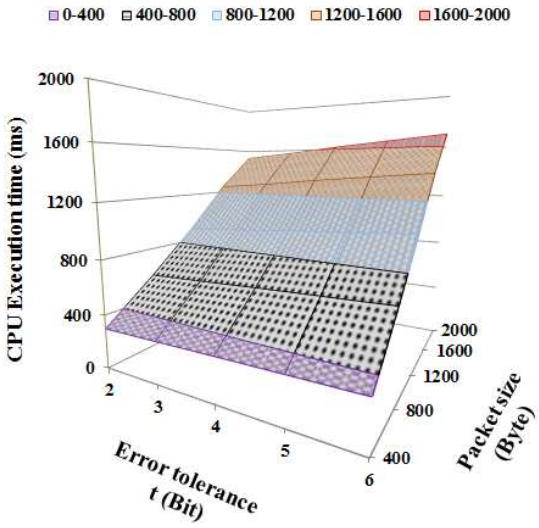

Fig. 7. Execution time of the CPU-based decoder.

In addition, greater partitioning of a packet results in a smaller segment size, which enables the GPU to achieve a faster execution time than the CPU.

Overall, the GPU outperforms the CPU in terms of time required to decode various packet sizes and error tolerance, yielding a tremendous improvement in execution time. We attribute this to the massively parallel design of SP, DE, and ME of the GPU. Detailed speedup information is summarized in Table 1; the maximum speedup gained by the GPU over the CPU is 99×.

Table 1. Speedup of GPU over CPU.

Packet size	t=2	t=3	t=4	t=5	t=6
M=400	13×	14×	16×	18×	21×
M=800	26×	27×	30×	35×	40×
M=1200	38×	40×	45×	51×	60×
M=1600	52×	54×	61×	70×	81×

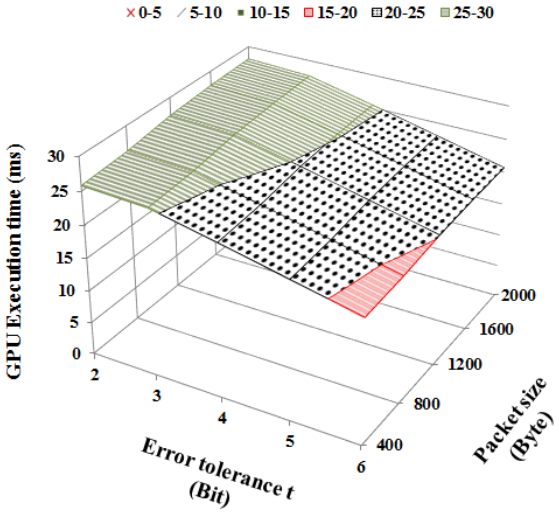

Fig. 8. Execution time of the GPU-based decoder.

V. Conclusions

In this paper, we proposed a computationally efficient GPU implementation of a Hamming code decoder for faster error recovery in data communication networks. We compared the performance of the proposed GPU approach with an equivalent sequential approach on a traditional CPU. The GPU-based implementation strongly outperformed the CPU-based sequential approach in terms of execution time, yielding a 99× speedup. These results indicate that the proposed GPU approach is suitable for application in time-sensitive and high-speed wired and wireless communication systems.

REFERENCES

[1] U. Ramacher, "Software-Defined Radio Prospects for Multistandard Mobile Phones," *IEEE Journal of Computer*, vol.40, no.10, pp.62-69, Oct. 2007.

[2] T. Yazdi, S.M.Sadegh, H. Cho, L. Dolecek, "Gallager B Decoder on Noisy Hardware," *IEEE Transactions on Communications*, vol.61, no.5, pp.1660-1673, May 2013.

[3] S. Gronroos, K. Nybom, and J. Bjorkqvist, "Complexity analysis of software defined DVB-T2 physical layer," *Journal of Analog Integrated Circuits and Signal Processing*, vol. 69, no. 2-3, pp. 131-142, Dec. 2011.

[4] A. Refaey, S. Roy, P. Fortier, "A new approach for FEC decoding based on the BP algorithm in LTE and WiMAX systems," in *Proc. 12th IEEE Canadian Workshop on Information Theory (CWIT)*, pp.9-14, 17-20 May 2011.

[5] E. Nicollet, "Standardizing Transceiver APIs for Software Defined and Cognitive Radio," *IEEE RFDesign Magazine*, vol. 47, no. 1, pp. 16-20, Feb. 2008.

[6] M. May, T. Ilnseher, N. Wehn, and W. Raab, "A 150Mbit/s 3GPP LTE turbo code decoder," in *Proc. ACM Conference on Design, Automation and Test in Europe (DATE '10)*, pp. 1420-1425, 2010.

[7] Y. Lin, H. Lee, M. Who, Y. Harel, S. Mahlke, T. Mudge, C. Chakrabarti, K. Flautner, "SODA: A High-Performance DSP Architecture for Software-Defined Radio," *IEEE journal of Microcomputer and Microprocessor Systems*, vol.27, no.1, pp.114-123, Jan.-Feb. 2007.

[8] J. Y. Park, K.S. Chung, "Parallel LDPC decoding using CUDA and OpenMP," *EURASIP Journal on Wireless Communications and Networking*, Springer, vol.11, no.1, pp.172-180, Nov. 2011.

[9] M. Palkovic, P. Raghavan, M. Li, A. Dejonghe, L. Van der Perre, F. Catthoor, "Future Software-Defined Radio Platforms and Mapping Flows," *IEEE Magazine of Signal Processing* , vol.27, no.2, pp.22-33, Mar. 2010.

[10] W. Tuttlebee, "The Software Defined Radio: Enabling Technologies," *John Wiley & Sons*, 2002.

[11] J. Kim, S. Hyeon, S. Choi, "Implementation of an SDR system using graphics processing unit," *IEEE Communications Magazine*, vol.48, no.3, pp.156-162, Mar. 2010.

[12] T. Beluch, F. Perget, J. Henaut, D. Dragomirescu, R. Plana, "Mostly Digital Wireless UltraWide Band Communication Architecture for Software Defined Radio," *IEEE Microwave Magazine*, vol.13, no.1, pp.132-138, Jan.-Feb. 2012.

[13] M. Wu, Y. Sun, S. Gupta, and J. R. Cavallaro, "Implementation of a High Throughput Soft MIMO Detector on GPU," *Journal of Signal Processing Systems*, vol. 64, no. 1, pp. 123-136, 2011.

[14] H. Lee, C. Chakrabarti, T. Mudge, "A Low-Power DSP for Wireless Communications," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol.18, no.9, pp.1310-1322, Sep. 2010.

[15] C. S. Lin, W. L. Liu, W. T. Yeh, L. W. Chang, W. M. W. Hwu, S. J. Chen, P. A. Hsiung, "A Tiling-Scheme Viterbi Decoder in Software Defined Radio for GPUs," in *Proc. 2011 7th International Conference on Wireless Communications, Networking and Mobile Computing (WiCOM)*, pp.1-4, 23-25 Sept. 2011.

[16] V. B. Alluri, J. R. Heath, M. Lhamon, "A New Multichannel, Coherent Amplitude Modulated, Time-Division Multiplexed, Software-Defined Radio Receiver Architecture, and Field-Programmabale-Gate-Array Technology Implementation," *IEEE Transactions on Signal Processing*, vol. 58, no. 10, pp.5369-5384, 2010.

[17] L. Hu, S. Nooshabadi, T. Mladenov, "Forward error correction with Raptor GF(2) and GF(256) codes on GPU," *IEEE Transactions on Consumer Electronics*, vol.59, no.1, pp.273-280, Feb. 2013.

[18] Y. Zhao, F.C.M. Lau, "Implementation of Decoders for LDPC Block Codes and LDPC Convolutional Codes Based on GPUs," *IEEE Transactions on Parallel and Distributed Systems*, vol.25, no.3, pp.663-672, Mar. 2014.

[19] W. Michael, S. Yang, G. Wang, and C. Joseph, "Implementation of a High Throughput 3GPP turbo decoder on GPU," *Journal of Signal Process System*, vol. 65, no. 1, pp.171-183, 2011.

[20] F. J. Martinez-Zaldivar, A. M. Vidal-Macia, A. Gonzalez, and V. Almenar, "Tridimensional block multiword LDPC decoding on GPUs," *Journal of Supercomputing*, vol. 58, no. 3, pp. 314-322, 2011.

[21] R. Li, J. Zhou, Y. Dou, S. Guo, Dan Zou, and S. Wang, "A multi-standard efficient column-layered LDPC decoder for Software Defined Radio on GPUs," in *Proc. IEEE 14th Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pp.724-728, 16-19 Jun. 2013.

[22] R. Li, Y. Dou, Y. Li, S. Wang, "A fully parallel truncated Viterbi decoder for Software Defined Radio on GPUs," in *Proc. 2013 IEEE Wireless Communications and Networking Conference (WCNC)*, pp.4305-4310, 7-10 Apr. 2013.

[23] W. Michael, S. Yang, G. Siddharth, and C. Joseph, "Implementation of a High Throughput Soft MIMO Detector on GPU," *Journal of Signal Process System*, vol. 64, no. 1, pp. 123-136, 2011.

[24] C. Ahn, J. Kim, J. Ju, J. Choi, B. Choi, and S. Choi, "Implementation of an SDR platform using GPU and its application to a 2x2 MIMO WiMAX system," *Journal of Analog Integrated Circuits and Signal Processing*, vol. 69, no. 2, pp. 107-117, 2011.

[25] R. W. Hamming, "Error detecting and error correcting codes," *The Bell System Technical Journal*, vol. 26, no. 2, pp. 147-160, 1950.

[26] J. Xu, K. Li, and G. Min, "Reliable and Energy-Efficient Multipath Communications in Underwater Sensor Networks," *IEEE Transactions on Parallel and Distributed Systems*, vol.23, no.7, pp.1326-1335, Jul. 2012.

[27] R.Ma, S. Cheng, "The Universality of Generalized Hamming Code for Multiple Sources," *IEEE Transactions on Communications*, vol.59, no.10, pp.2641-2647, Oct. 2011 .

[28] S. Islam, J. Kim, "Accelerating Extended Hamming Code Decoders on Graphic Processing Units for High Speed Communication," IEICE Trans. on Communications, Vol.E97-B, No.5, pp.1050-1058, May 2014.

[29] A.G. Amat, C. A. Nour, and C. Douillard, "Serially concatenated continuous phase modulation for satellite communications," *IEEE Transactions on Wireless Communications*, vol.8, no.6, pp.3260-3269, Jun. 2009.

[30] C.Argyrides, R. R. Ferreira, C.A. Lisboa, and L. Carro, "Decimal Hamming: A Software-Implemented Technique to Cope with Soft Errors," in *Proc. 2011 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems*, pp.11-17, 3-5 Oct. 2011.

[31] R.Ma, and S. Cheng, "Hamming coding for Multiple Sources," in *Proc. 2010 IEEE International Symposium on Information Theory Proceedings ISIT)*, pp.171-175, 13-18 Jun. 2010.

Md Shohidul Islam received a B.S. degree in computer science and engineering from Rajshahi University of Engineering and Technology, Rajshahi, Bangladesh, in 2007, and an M.S. in computer engineering from the University of Ulsan, Ulsan, South Korea, in 2014. His current research interests include high-speed error coding, GPU computing, software defined radio (SDR), parallel computing, and the performance evaluation of many-core processors for application-specific SoC design.

Cheol-Hong Kim received a BS, a MS, and a PhD in computer engineering from Seoul National University, Seoul, Korea, in 1998 2000, and 2006, respectively. He is an Associate Professor of Electronics and Computer Engineering at Chonnam National University, Gwangju, Korea. His research interests include multicore architecture and embedded systems.

Jong-Myon Kim received a B.S. degree in electrical engineering from Myongji University, Yongin, South Korea, in 1995, an M.S. in electrical and computer engineering from the University of Florida, Gainesville, in 2000, and a Ph.D. in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, in 2005. He is an Associate Professor of Electrical Engineering at the University of Ulsan, Ulsan, South Korea. His research interests include GPU computing, multimedia specific processor architecture, parallel processing, and embedded systems. He is a member of IEEE and IEICE.